arXiv:1406.5634v1 [cs.NI] 21 Jun 2014

A Framework to Quantify the Benefits of Network Functions Virtualization in Cellular Networks

Zafar A. Qazi†, Vyas Sekar⋆, Samir R. Das†
† Stony Brook University ⋆ Carnegie Mellon

Abstract

Network functions virtualization (NFV) is an appealing vision that promises to dramatically reduce capital and operating expenses for cellular providers. However, existing efforts in this space leave open broad issues about how NFV deployments should be instantiated or how they should be provisioned. In this paper, we present an initial attempt at a framework that will help network operators systematically evaluate the potential benefits that different points in the NFV design space can offer.

1 Introduction

Cellular networks today incur high capital costs in deploying a broad range of expensive and inflexible hardware appliances which include both cellular-specific functions such as Serving and Packet Data Networks Gateway (S/P-GW), IP Multimedia System (IMS) elements like Call Session Control Functions (CSCFs) [3], as well as more traditional network appliances (e.g., NATs, firewalls, proxies) [18]. Furthemore, they incur high management complexity on many fronts: diverse protocols and standards (e.g., 3G, 4G), multiple types of services (e.g., video, voice, messaging), and fine-grained policy requirements (e.g., per-user accounting of data and voice calls and specialized video services).

In conjunction, these effects have led to a state where cellular providers face an uphill battle with the trend in the gap between revenues and their capital and operating expenses being quite unfavorable [15]. Furthermore, current deployments are inflexible on several accounts: they cannot react to changing demands and policies and the timescales of innovation are hindered by vendor support.

This trend is likely unsustainable in the longer term and has motivated the case for *network functions virtualization* (NFV) [4]. The motivation in NFV is to bring the benefits that cloud computing has provided for the IT industry to network operators: accelerating the pace of innovation by reducing the cycles to deploy new equipment, economies of scale provided by commodity hardware, resource multiplexing via virtualization, dynamic provisioning and elastic scaling, and the ability to experiment with new services without significant upfront costs [4](see also Figure 1). The high-level idea is that the various *network functions* (NFs) for cellular-, IP-, and application-specific services can be replaced by *virtualized* applications on commodity hardware platforms.

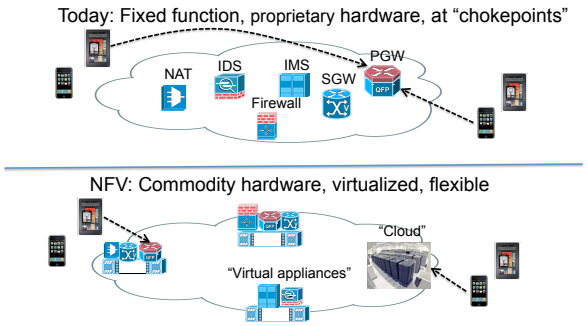

Figure 1: The current fixed and proprietary implementation of network functions *vs* the NFV vision of elastic, cost effective, mix-match and potentially hybrid deployment of network functions.

As highlighted in an early whitepaper [4], this vision builds upon, and is complementary, to existing work in the software-defined networking (SDN). The main differences are that: (a) it broadens the scope of the data plane functions beyond OpenFlow-enabled switches; and (b) it focuses more on the carriers and the services they would like to offer to their customers.

While the above long-term vision is appealing, it also leaves open several aspects of the design space that operators would need to address in practice w.r.t.:

- the type of *platforms* (e.g., pure cloud, or fixed-but-flexible hardware infrastructure);
- the type of *provisioning* (i.e., where to place new flexible hardware or datacenters); and
- *demand distribution* (i.e., how to route user demands to different function instances to meet logical policy requirements).

Given this broad design space, there are several possible NFV instantiations by combining choices across the individual platform, provisioning, and distribution dimensions. For instance, we can imagine a consolidated deployment where the provider has a small number of datacenters at which the network functions (NFs) can be run in a dynamic, elastic manner. At the other extreme, we can imagine current deployments with specialized NF hardware. We can also envision "nano-datacenter" like models with flexible NF hardware distributed throughout the network [17]. These deployments will naturally have different provisioning, operational, and performance characteristics.

Given this diverse and broad design space, network

operators will need systematic decision systems to help them evaluate the cost-benefit tradeoffs of different points in the design space. We highlight some motivating scenarios in §2. Furthermore, even before embarking on rearchitecting their network infrastructure [2], operators need to first *quantify* the potential CAPEX and OPEX benefits that specific NFV strategies might offer.

This paper is a first attempt to shed light on these issues. To this end, we cast the NFV deployment problem as a systematic optimization framework (§4). Our framework is general enough to capture different points in the design space and also consider hybrid NFV different deployments (e.g., some combination of pure cloud and fixed hardware). In order to estimate the potential benefits of different NFV strategies, the operators provide as input historical traffic demands and policy based service chaining requirements for different traffic patterns (§3). Our framework will then output guidelines on the optimal provisioning strategy and the cost benefits it offers. Operators can use such a framework for "what-if" analysis to evaluate the cost-benefit tradeoffs of different deployment strategies.

As illustrative examples we show how operators can use our framework to evaluate the benefits of different NFV designs (§5). For instance, we observe that using flexible hardware minimizes the deployment cost in many scenarios. We also conduct a sensitivity analysis to evaluate the effects of changing different input parameters on the optimal deployment strategy.

2 Motivation

We begin by outlining the NFV design space and then use motivating scenarios to highlight how the optimal NFV strategy depends on the workload and cost factors.

2.1 Design Space of Cellular NFV

We identify three key dimensions for the design space:

- *Platform type:* Network functions (NFs) can be realized in many ways. Today, each NF is a dedicated appliance (*Single*) providing a specialized capability. Going forward, one can imagine a flexible commodity hardware (*FlexHW*) that can be repurposed to run different types of NFs on demand [11, 14, 6]. Going one step further, we can imagine that the functions are themselves outsourced to a *Cloud* service that can elastically scale resources for different NFs.[1]
- *Provisioning and placement:* A key operational decision is deciding how and where to provision NF platforms. At one extreme, the provider can choose a single *Cloud* location. At the other extreme we can envision a nano-datacenter model where every cell base station has an associated mini-*Cloud* (e.g., [17]). We can also consider simple hybrids where each location has pre-provisioned *Single* and *FlexHW* boxes and we have a few *Cloud* locations.
- *Scaling and distribution strategies:* Given a specific provisioning/placement strategy, another aspect of the design space is how the available hardware resources (possibly elastic) are used to serve the (varying) offered load. Again, we can envision several possible strategies here: optimal load balancing, or routing to nearest available instance of a specific NF, or offloading to the cloud beyond a threshold value of load.

[1]Industry reports use the terms *FlexHW* and *Cloud* rather loosely. They are however quite distinct in their cost-performance tradeoffs and thus one of our goals in formalizing the NFV problem space is to crystallize these loose characterizations of NFV instantiations.

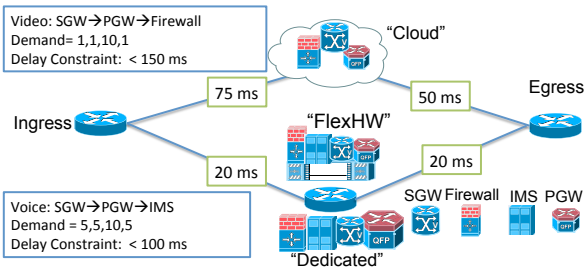

Figure 2: Example to motivate the different design tradeoffs in provisioning.

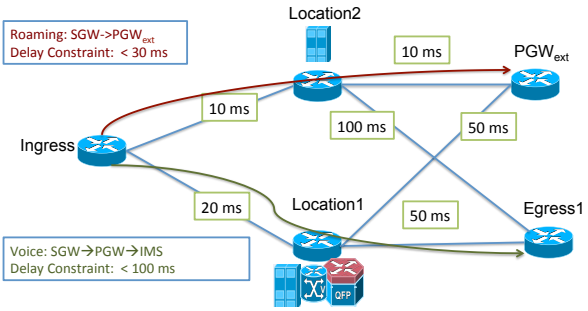

Figure 3: Example to illustrate the different design tradeoffs in functional placement and routing.

2.2 Motivating Scenarios

We use the simple scenarios in Figure 2 and Figure 3 to illustrate how different points in the NFV design space may be optimal depending on the performance constraints, traffic patterns, policy chains and CAPEX and OPEX.

Optimizing provisioning cost: Consider the single traffic class, *Video*, with traffic entering at "Ingress" and exiting at "Egress" with the *service chain* SGW-PGW-FIREWALL. Suppose the traffic volume for this traffic class across four time epochs is $\{1,1,10,1\}$. Assume that the (fixed) cost per unit capacity for fixed in-network provisioning is 20. On the other hand, if we use *Cloud* to dynamically provision the needed resource, assume the amortized cost per unit capacity provisioned is 10. With fixed pre-provisioning we need to allocate for peak

traffic and thus a capacity of 10 units costing 200. But when outsourcing the processing to the cloud, the cost is only 130. Thus, we decide to provision the entire service chain for *Video* in the *Cloud*.

Performance constraint: Now, assume that we add another traffic class, *Voice*, with the *service chain* SGW-PGW-IMS. Assume that *Voice* has a performance constraint that the average latency should be less than 100 ms. Then, we cannot use the *Cloud* for processing traffic belonging to this class. Assume also that the traffic volume for *Voice* across the four time epochs is {5,5,10,5}. Now, a hybrid solution where we provision a capacity of 10 units on fixed hardware and 10 units on the *Cloud* dynamically only in epoch 3 will be the most cost effective. This solution has cost 300, while provisioning entirely on the *Cloud* costs 380 (if there are no performance constraints) and provisioning entirely on a fixed hardware costs 400.

Functional placement and routing: To illustrate placement/routing issues consider the network in Figure 3. We have two traffic classes *Roaming* and *Voice*. The traffic belonging to *Roaming* requires processing SGW-PGW_{ext}, where PGW_{ext} is the PGW of the external network to which the traffic belongs. On the other hand, *Voice* needs processing using the service chain SGW-PGW-IMS. Assume that the carrier normally prefers to provision all resources in "Location 1" due to cost or management issues. However, when the *Roaming* traffic enters the network, it must move the related SGW processing to "Location 2" as going via "Location 1" violates the delay constraint for *Roaming*. Assume also after a network upgrade the delay on the Location 1-PGW_{ext} link falls to 5 ms. Then, the SGW processing can move back to Location 1. Moving the SGW processing back and forth depending on traffic and network changes could be facilitated via use of *FlexHW*.

The above scenarios highlights the tradeoffs and considerations operators need to make. It is important for operators to be able to analyze these tradeoffs before rolling out new NFV deployments. However, given the size and complexity of modern cellular deployments, it may be impractical for operators to manually evaluate the entire space of design options. Thus, our goal is to develop a decision support framework that allows cellular operators to systematically explore different design options before deploying new hardware. For instance, they would like to specify policy requirements (e.g., service chains for different user classes) and performance constraints (e.g., load, congestion, or latency) and use such a decision support system to choose the right mix of platform, provisioning, and distribution strategies from the broader NFV design space highlighted above.

3 Inputs and Requirements

We begin by describing the requirements and inputs that operators need to provide to our framework. This data can be obtained from their network logs, policy configurations, and vendor-specific benchmarks.

Traffic patterns: Cellular traffic is divided into different logical *classes* based on different user/customer demands. For example, the classes may capture *regular users* vs. *roaming customers* vs. *machine-to-machine* (M2M) traffic, with different requirements. We assume that the operator has historical demand patterns, with $|T_{c,e}|$ representing the volume of traffic for class c observed in epoch e.

Processing requirements: Each class c is associated with a *policy service chain* or a sequence of *network functions* (NF) that process traffic in c. Let $SC_c = NF_1 \prec NF_2 \prec \ldots$ denote the service chain for class c. These NFs could span cellular-, IP-, and application-level processing. Let $FP_{c,m,t}$ denote the processing cost (e.g., CPU usage) per-packet in class c for running a NF NF_m on a specific *type* of NF platform t. For example, the per-packet CPU usage may differ across virtualized vs. non-virtualized deployments.

Performance constraints: Performance constraints specify that traffic in class c should have some pre-specified performance PC_c [3]. As a simple starting point, we consider the end-to-end latency for each class.

Cost factors: The provisioning costs associated with rolling out the NF platforms may depend on the platform t (i.e., *FlexHW* or *Cloud*) and the location l (e.g., power, cooling costs). We capture these costs as follows. First, we assume that there is a *fixed* cost of deploying an instance of type t at location l, $Fixed_{l,t}$; e.g., this captures administrative and labor costs in rolling out new deployments and the operational costs. For deployments like *Cloud* this may be zero as they may use a pay-as-you-go model. Second, there is a *hardware* cost, $Var_{l,t}$, depending on the amount of resource provisioned for this instance; e.g., based on the number of CPU cores or memory on the hardware. Third, we have a *elastic* factor depending on the actual resources used, $Elas_{l,t}$; e.g., this can be linear in the amount of resources used for *Cloud* and zero for the others.

4 Provisioning Model

In this section, we describe a formal optimization framework that captures the cost of provisioning the network to meet the time-varying processing and performance requirements, given the policy constraints, platform costs, traffic demands, and network specifications. Network operators can use this framework to (a) systematically quantify the potential benefits of NFV in their deployments and (b) estimate the relative benefits offered by

different points in the NFV design space.

4.1 Control Variables

As a starting point, we present a model where the operator is considering a set of possible platform instances p to deploy. As discussed in the previous section, each p can be instantiated in many ways (e.g., *Single* vs. *Cloud*) with varying cost-performance characteristics. We use $t(p)$ and $l(p)$ to denote the type (e.g., *Cloud* or *Single*) and location of a specific NF platform instance p. Now, there are two main types of control variables that we need to capture:

- *Provisioning:* The first decision is a binary decision if we want to deploy an instance at a specific location. This is captured by a $\{0,1\}$ variable $Active_p$. If we choose to deploy, then we also need to decide how much hardware resource to provision, captured by the variable Res_p. For some platform types, we also have an elastic option (e.g., *FlexHW* or *Cloud*), in this case we also use dynamic provisioning decision variables, $Res_{p,e}$.
- *Load distribution:* Given a provisioning strategy, we need to meet the processing requirements and distribute the load across the various p instances. In general, each class may have different chains and the required NFs can be instantiated at *any* set of instances capable of running these NFs. We can also flexibly route traffic to balance the network and platform loads [13]. To capture these considerations, we introduce *flow variables*, $f_{c,e,p_m,p'_{m'}}$, that represents the fraction of traffic in class c in epoch e, routed from a NF instance p_m to another NF instance $p'_{m'}$ (similar to Figure 4). Note that we can flexibly capture different routing strategies by scoping these flow distribution variables differently. (Some of these variables will not appear if m or n do not appear in SC_c.)

Objective function: Our objective is to minimize the total provisioning cost. This has three components: (1) fixed costs of instantiating platforms; (2) hardware costs for each "active" platform; and (3) the dynamically provisioned compute resources per epoch. Eq (1) shows this total cost in terms of the $Active_p$, Res_p, and $Res_{p,e}$ control variables:

$$\sum_p Fixed_{l(p),t(p)} \times Active_p + Var_{l(p),t(p)} \times Res_p + \sum_{p,e} Elas_{l(p),t(p)} \times Res_{p,e} \quad (1)$$

As discussed earlier, these factors depend on the type of platform; e.g., *Cloud* may have $Fixed_{l(p),t(p)} = Var_{l(p),t(p)} = 0$, but have $Elas_{l(p),t(p)} > 0$, while other models have $Elas_{l(p),t(p)} = 0$.

4.2 Formulation

Next, we describe how we capture the various processing and provisioning constraints.

Resources provisioned: First, we need to capture the amount of resources provisioned. We begin by capturing the total compute load on a NF instance p_m on the platform p during epoch e in Eq (2):

$$\forall e,p,m : LoadPerNF_{p_m,e} = \sum_c \sum_{\substack{m' \text{ s.t } m' \in SC_c \\ \text{and } m' \in p'}} FP_{c,m,t} \times f_{c,e,p_m,p'_{m'}} \times |T_{c,e}| \quad (2)$$

Then, in Eq (3), the total load on a platform p in an epoch e is simply the sum over all NF functions that p can support:

$$\forall e,p : Load_{p,e} = \sum_{m \in p} LoadPerNF_{p_m,e} \quad (3)$$

Now, our provisioning strategy must ensure that each NF platform has sufficient resources to cover the processing requirements per epoch and across all epochs. Thus, we have Eq (4) and Eq (5):

$$\forall e,p : Load_{p,e} = Res_{p,e} \quad (4)$$

$$\forall e,p : Load_{p,e} \leq Res_p \quad (5)$$

In addition, and depending on other capacity constraints (e.g., space, power, or available hardware configurations), we may also have upper bounds on the total resources per-platform at each location:

$$\forall p : Res_p \leq Cap_p \quad (6)$$

Fixed costs: Next, we need to model the fixed costs associated with the above provisioning strategy. These fixed costs are incurred if the platform is being used in *at least* one of the epochs with non-zero resources. Thus, we have the following relationship between the binary $Active_p$ variables and the Res_p variables:

$$\forall p : Res_p \leq Cap_p \times Active_p \quad (7)$$

Modeling traffic distribution: The above equations model the provisioning aspects, but do not capture how the traffic processing is *distributed* across the platforms. In other words, we need to model how the $f_{c,e,p_m,p'_{m'}}$ variables are quantified. There are two key things we need to capture here. First, we need to model the *coverage* constraint that for each c, the desired service chain SC_c has been assigned to some set of platform instances. Second, we also need to ensure that the service chain is correctly applied in the intended *sequence*. Let $SC_c[j]$ denote the j^{th} NF in the chain SC_c. Note that each $SC_c[j]$ may be realized using several candidate platform instances.

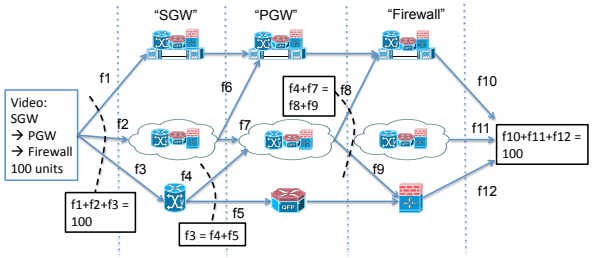

Figure 4: Example to illustrate how flow conservation is modeled.

We model this using two sets of constraints. First, in Eq (8) we ensure that the entire fraction of traffic is routed to the first hops.

$$\forall c, e: \sum_{p_m : m = SC_c[1]} f_{c,e,p_m} = 1 \quad (8)$$

Second, we model *flow conservation* constraints that ensures that the traffic incoming into one "stage" in the service chain is routed to the next "stage" in the desired sequence(e.g, see Figure 4).[2] Then, we have:

$$\forall p_m, c, e \text{ s.t } m = SC_c[j] \ \& \ j > 1: \sum_{p'_{m'} : m' = SC_c[j-1]} f_{c,e,p_m,p'_{m'}} = \sum_{p'_{m'} = SC_c[j+1]} f_{c,e,p_m,p'_{m'}} \quad (9)$$

Performance bounds: Now, given the flow distribution variables $f_{c,e,p_m,p'_{m'}}$, we can also model the network-level performance that traffic for each class perceives. As a simple starting point, we model the *average latency* and ensure that this is less the given threshold PC_c. If $Lat_{p_m,p'_{m'}}$ is the typical network latency on the path from p_m to $p'_{m'}$, then we can capture the performance bound as shown below:

$$\forall c, e: \sum_{j=1}^{|SC_c|-1} \sum_{\substack{p_m, p'_{m'} \text{ s.t} \\ m = SC_c[j] \\ m' = SC_c[j+1]}} f_{c,e,p_m,p'_{m'}} \times Lat_{p_m,p'_{m'}} \leq PC_c \quad (10)$$

5 Example Use Cases

Next, we highlight some illustrative use cases to validate how operators can use our framework to evaluate NFV design tradeoffs.

Setup: Due to lack of publicly available information on cellular network topologies, we use the PoP-level Internet2/Abilene topology.[3] We currently use four traffic classes, each with three different NFs. We assume there are a total of 8 different NFs [3]. We use a gravity model based on city populations as a baseline traffic demand and simple randomized variability models. We model the different provisioning problems as an integer linear program (ILP) and use CPLEX to solve the problem.

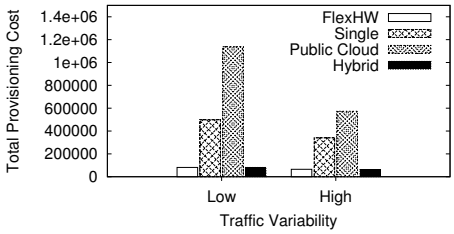

Figure 5: Total provisioning cost for different NFV models.

We instantiate the different platform types as follows:

- *FlexHW* assumes a commodity server is being used with standard virtualization technologies to run different functions on a single platform;
- *Single* assumes specialized hardware for running single functions;
- *Public Cloud* outsource the processing to a public cloud provider; and
- A *Hybrid* deployment model which allows full flexibility to use any combination of the above three deployment models.

Again, given the absence of accurate cost numbers for NFV platforms, we obtain ballpark numbers for the different costs by the following strategy. Since the different platforms have fundamentally different cost/service models, we normalize the costs by computing the dollar cost in provisioning/running the platform for unit traffic; e.g., dollars-per-Mbps. First, for *Public Cloud*, we consider bandwidth costs in Amazon EC2[4], and we compute the normalized cost assuming a monthly transfer volume of 500 TB. Second, for the *FlexHW* hardware we assume a typical commodity server of price $2,500 as representative for *FlexHW*. Third, for the *Single* devices, we use numbers from published work and assume a specialized device at 20 Gbps capacity costs roughly $80,000 [12] as representative for *Single* devices. Finally, for the setup and operational cost, we use a common industry rule of thumb and model it to be twice the equipment cost [1].

Comparing different design options: To illustrate this point, Figure 5 shows that with the given costs a deployment strategy using only *FlexHW* has the same provisioning cost as the *Hybrid* when the traffic variability is low and very close to *Hybrid* when the variability is high.

Next, we show the impact of varying different input parameters like cloud cost, setup and operational cost,

[2] In this paper, we are not we are not mandating a specific data plane implementation. We could use wildcard rules [19] or tunnels [13].

[3] We have also experimented with other ISP topologies from RocketFuel.

[4] The CPU, memory costs were much lower than bandwidth costs

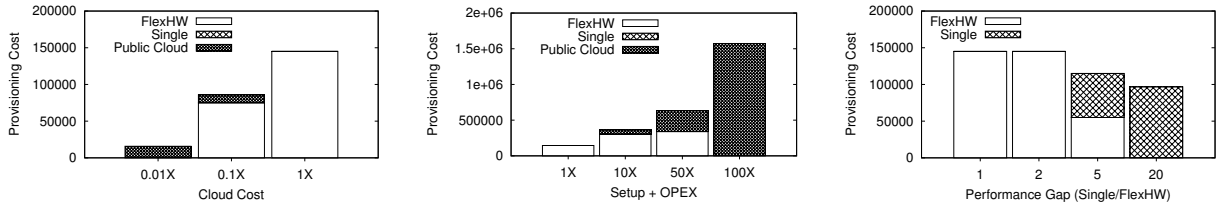

(a) Varying cloud cost (b) Varying setup+OPEX cost (c) Varying performance gap

Figure 6: Sensitivity of the optimal strategy w.r.t. cost and performance considerations.

and performance of each platform type. We consider a *Hybrid* deployment model for these experiments with a random traffic matrix. Figure 6a shows that as the *Cloud* cost decreases, *Cloud* becomes a more viable option for processing of network functions, with a *Cloud* cost of *0.01X*, resulting in a *Cloud* only optimal strategy. Figure 6b shows that as the setup and operational cost increases, there is less incentive in pre-provisioning resources inside the network. Figure 6c shows that as the performance gap between virtual appliances and specialized hardware increases, it maybe cost effective to use a combination of these two types of platforms. Note, for Figure 6c we only consider a *Hybrid* model consisting of *FlexHW* and *Single* platforms.

6 Related Work

To the best of our knowledge, there have been no systematic frameworks to characterize the NFV design space and provide tools for operators to explore "what-if" deployment scenarios. We discuss three complementary threads of related work below.

Cellular SDN and NFV: Industry reports and actual measurements have confirmed that there are a large number of complex function required in the cellular core [5, 18]. Today, these networks are inflexible and expensive to provision/maintain. This has motivated the case for network functions virtualization [4]. SoftCell [10] proposes an SDN based cellular core that can support very fined-grained policies. Our work is complementary as we focus on provisioning and placement decisions.

Middleboxes: Related work has focused on "middlebox" service chaining and load balancing [13, 14, 20]. Other work has also suggested NFV-like ideas for traditional middleboxes [16, 9]. In our work, we consider the provisioning and placement in addition to these requirements. Moreover, we present a general optimization framework that captures all three key aspects of the NFV design space: type of platforms, provisioning, and demand distribution.

Cloud provisioning: Prior work has tried to formulate the potential cost savings in cloud computing [8, 7]. As the NFV vision begins to gain momentum, we believe that there is a critical need for frameworks like ours to shed light on the quantitative benefits that different NFV strategies may offer.

7 Conclusions and Future Work

Our framework is only a first step and we identify several natural directions for future work and extensions. First, our framework focuses largely on how NFV can simplify existing deployments. We are actively working with operators to understand how the dynamicity and flexibility offered by NFV can enable new deployment/service opportunities. Second, we are also working with operators to obtain more realistic datasets and policy requirements to quantify the benefits of NFV in practice. Finally, we can also extend our model to incorporate other kinds of policy considerations; e.g., to make sure that roaming users are not co-located on the same physical hardware with other users.

8 References

[1] A Simple Model for Determining True Total Cost of Ownership for Data Centers. `http://tinyurl.com/kznlhn2`.

[2] AT&T Domain 2.0 Vision White Paper. `http://tinyurl.com/p4uv3s3`.

[3] LTE Design and Deployment Strategies. `http://tinyurl.com/lj2erpg`.

[4] Network Functions Virtualisation. `http://portal.etsi.org/nfv/nfv_white_paper.pdf`.

[5] The LTE Network Architecture. `http://tinyurl.com/negszts`.

[6] J. Anderson et al. xOMB: Extensible Open Middleboxes with Commodity Servers. In *Proc. ANCS*, 2012.

[7] M. Armbrust et al. Above the clouds: A berkeley view of cloud computing. Tech Report No. UCB/EECS-2009-28.

[8] Y. Chen and R. Sion. To cloud or not to cloud? musings on costs and viability. In *Proc. SOCC*, 2011.

[9] A. Gember et al. Stratos: A Network-Aware Orchestration Layer for Middleboxes in the Cloud. In *Technical Report arXiv:1305.0209*, 2013.

[10] X. Jin et al. SoftCell: Scalable and Flexible Core Network Architecture. In *CoNEXT*, 2013.

[11] J. Martins et al. Enabling fast, dynamic network processing with clickOS. In *Proc. HotSDN*, 2013.

[12] P. Patel et al. Ananta: Cloud Scale Load Balancing. In *Proc. SIGCOMM*, 2013.

[13] Z. Qazi et al. SIMPLE-fying Middlebox Policy Enforcement Using SDN. In *Proc. SIGCOMM*, 2013.

[14] V. Sekar et al. Design and Implementation of a Consolidated Middlebox Architecture. In *Proc. NSDI*, 2012.

[15] S.Elby. Carrier Vision of SDN and future applications to achieve a more agile mobile businesss. Keynote at the OpenFlow World Congress, 2012.

[16] J. Sherry et al. Making Middleboxes Someone Else's Problem: Network Processing as a Cloud Service. In *Proc. SIGCOMM*, 2012.

[17] V. Valancius et al. Greening the Internet with Nano Data Centers. In *Proc. CoNext*, 2009.

[18] Z. Wang et al. An Untold Story of Middleboxes in Cellular Networks. In *Proc. SIGCOMM*, 2011.

[19] M. Yu, J. Rexford, M. J. Freedman, and J. Wang. Scalable Flow-Based Networking with DIFANE. In *Proc. SIGCOMM*, 2010.

[20] Y. Zhang et al. StEERING: A Software-Defined Network for Inline Service Chaining. In *Proc. ICNP*, 2013.